



US009385994B2

(12) **United States Patent**  
**Ke et al.**

(10) **Patent No.:** **US 9,385,994 B2**  
(45) **Date of Patent:** **\*Jul. 5, 2016**

(54) **NETWORK SECURITY DEVICE**

(71) Applicant: **Juniper Networks, Inc.**, Sunnyvale, CA (US)

(72) Inventors: **Yan Ke**, San Jose, CA (US); **Yuming Mao**, Milpitas, CA (US); **Jian Tong**, San Jose, CA (US); **Guangsong Huang**, Sunnyvale, CA (US)

(73) Assignee: **Juniper Networks, Inc.**, Sunnyvale, CA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 29 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **14/169,710**

(22) Filed: **Jan. 31, 2014**

(65) **Prior Publication Data**

US 2014/0215598 A1 Jul. 31, 2014

**Related U.S. Application Data**

(63) Continuation of application No. 13/302,808, filed on Nov. 22, 2011, now Pat. No. 8,654,779, which is a continuation of application No. 12/551,034, filed on Aug. 31, 2009, now Pat. No. 8,068,487, which is a

(Continued)

(51) **Int. Cl.**  
**H04J 3/24** (2006.01)  
**H04L 29/06** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **H04L 63/0236** (2013.01); **H04L 63/0272** (2013.01); **H04L 63/1416** (2013.01)

(58) **Field of Classification Search**

CPC . H04L 12/56; H04L 2012/56; H04L 12/6418; H04L 65/102; H04L 29/06176; H04L 63/0236; H04L 63/0272; H04L 63/1416; H04Q 11/04  
USPC ..... 370/389, 353-357, 396, 473, 471  
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

5,781,550 A 7/1998 Templin et al.  
5,842,040 A 11/1998 Hughes et al.

(Continued)

**OTHER PUBLICATIONS**

Julkunen et al., "Enhance Network Security with Dynamic Packet Filter", IEEE (1998), pp. 268-275.

(Continued)

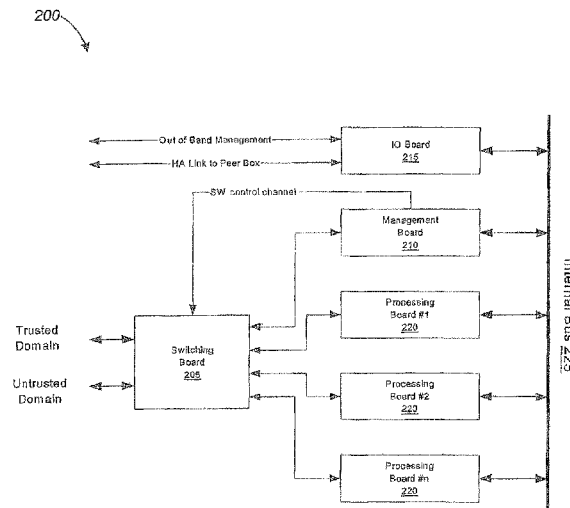
*Primary Examiner* — Brenda H Pham

(74) *Attorney, Agent, or Firm* — Harrity & Harrity, LLP

(57) **ABSTRACT**

Methods and apparatus, including computer program products, implementing and using techniques for processing a data packet. An input port receives a data packet, a switching board classifies the data packet, determines whether the data packet should be accepted, and switches the data packet to a management board if the data packet is a first data packet in a session, and to a processing board if the data packet is not a first data packet in a session. A management board receives a data packet from the switching board, examines the data packet and forwards the data packet to one of the processing boards. One or more processing boards receives non-first data packets from the switching board and data packets from the management board and processes the data packets. A firewall and a secure gateway with firewall and virtual private network functionality for processing a data packet are also described.

**20 Claims, 6 Drawing Sheets**



**Related U.S. Application Data**

continuation of application No. 11/428,235, filed on Jun. 30, 2006, now Pat. No. 7,602,775, which is a continuation of application No. 10/112,924, filed on Mar. 28, 2002, now Pat. No. 7,095,716.

- (60) Provisional application No. 60/280,684, filed on Mar. 30, 2001.

**References Cited****U.S. PATENT DOCUMENTS**

5,909,686 A 6/1999 Muller et al.  
 6,049,528 A 4/2000 Hendel et al.  
 6,088,356 A 7/2000 Hendel et al.  
 6,098,172 A 8/2000 Coss et al.  
 6,141,749 A 10/2000 Coss et al.  
 6,154,775 A 11/2000 Coss et al.  
 6,170,012 B1 1/2001 Coss et al.  
 6,252,878 B1 6/2001 Locklear et al.  
 6,263,368 B1 7/2001 Martin  
 6,591,303 B1 7/2003 Hendel et al.  
 6,606,315 B1 8/2003 Albert et al.  
 6,633,560 B1 10/2003 Albert et al.  
 6,650,641 B1 11/2003 Albert et al.  
 6,701,437 B1 \* 3/2004 Hoke et al. .... 726/15  
 6,704,278 B1 3/2004 Albert et al.  
 6,735,169 B1 5/2004 Albert et al.

6,742,045 B1 5/2004 Albert et al.  
 6,775,692 B1 8/2004 Albert et al.  
 6,798,788 B1 9/2004 Viswanath et al.  
 6,928,482 B1 8/2005 Ben Nun et al.  
 6,950,434 B1 9/2005 Viswanath et al.  
 7,023,863 B1 4/2006 Naudus et al.  
 7,032,037 B2 4/2006 Garnett et al.  
 7,042,870 B1 5/2006 Albert et al.  
 7,051,066 B1 5/2006 Albert et al.  
 7,095,716 B1 \* 8/2006 Ke ..... H04L 63/0272  
 370/230

7,143,438 B1 11/2006 Coss et al.  
 7,346,686 B2 3/2008 Albert et al.  
 7,535,907 B2 5/2009 Hussain et al.  
 7,602,775 B1 \* 10/2009 Ke ..... H04L 63/0272  
 370/386

7,643,481 B2 1/2010 Kadambi et al.  
 7,778,254 B2 8/2010 Kadambi et al.  
 7,895,431 B2 2/2011 Bouchard et al.  
 8,023,413 B2 9/2011 Kadambi et al.  
 8,068,487 B1 \* 11/2011 Ke ..... H04L 63/0272  
 370/389  
 8,654,779 B1 \* 2/2014 Ke et al. .... 370/401

**OTHER PUBLICATIONS**

Sharp et al., "Starburst: Building Next-Generation Internet Devices", Bell Labs Technical Journal 6(2), pp. 6-17 (2002).

\* cited by examiner

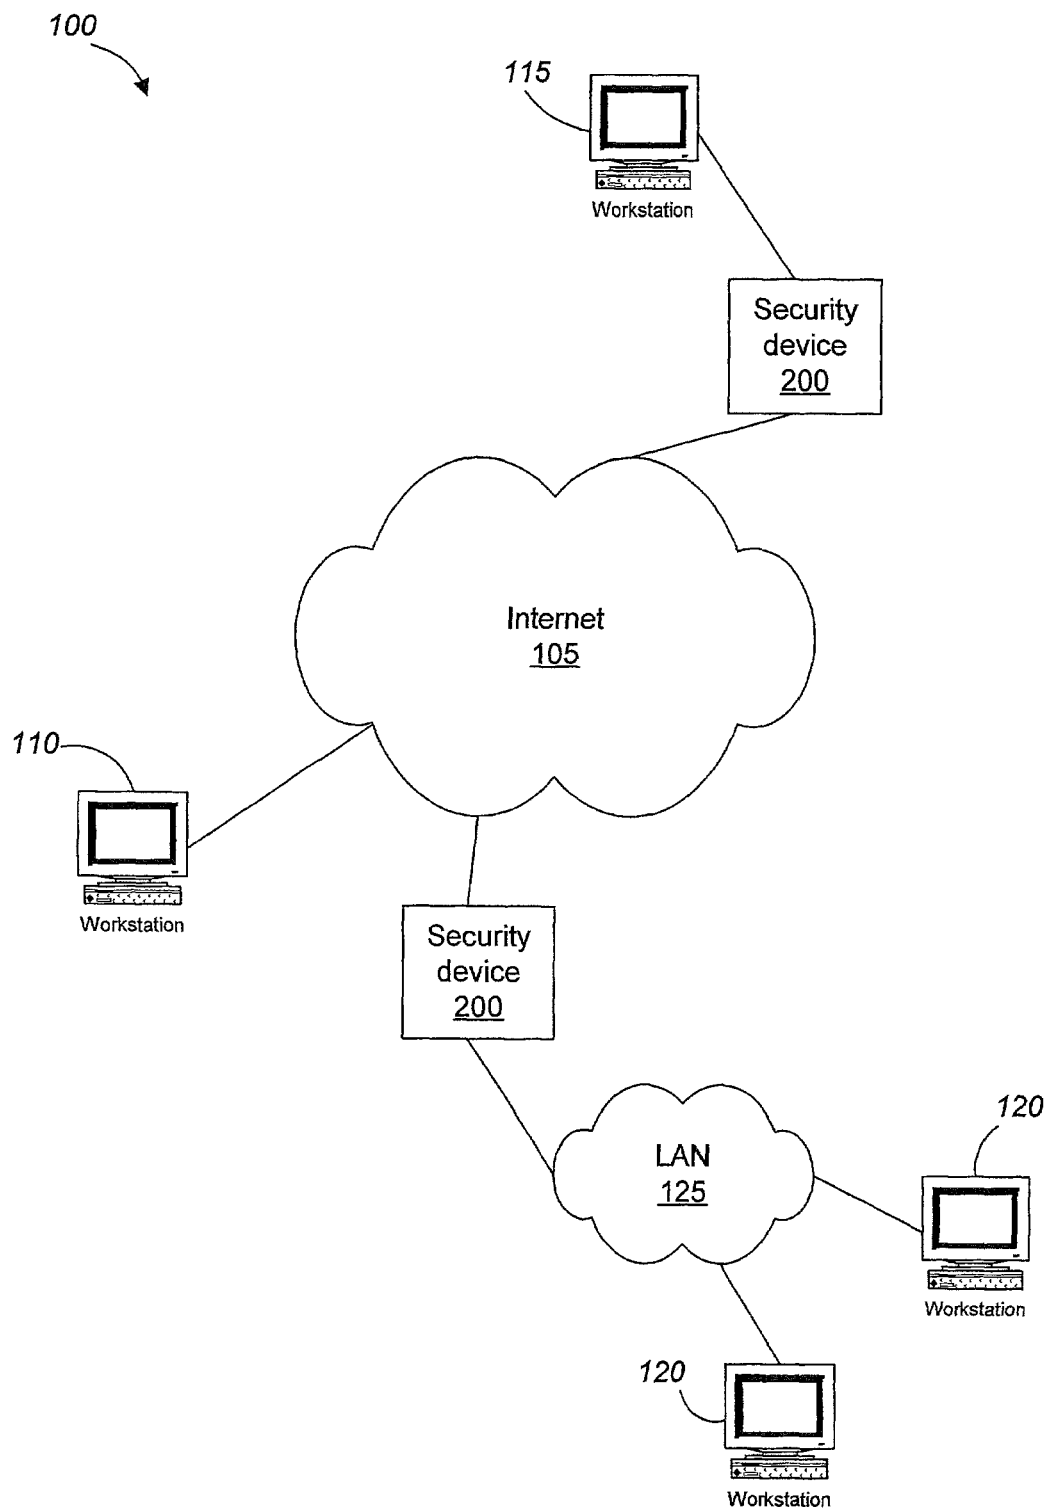


FIG. 1

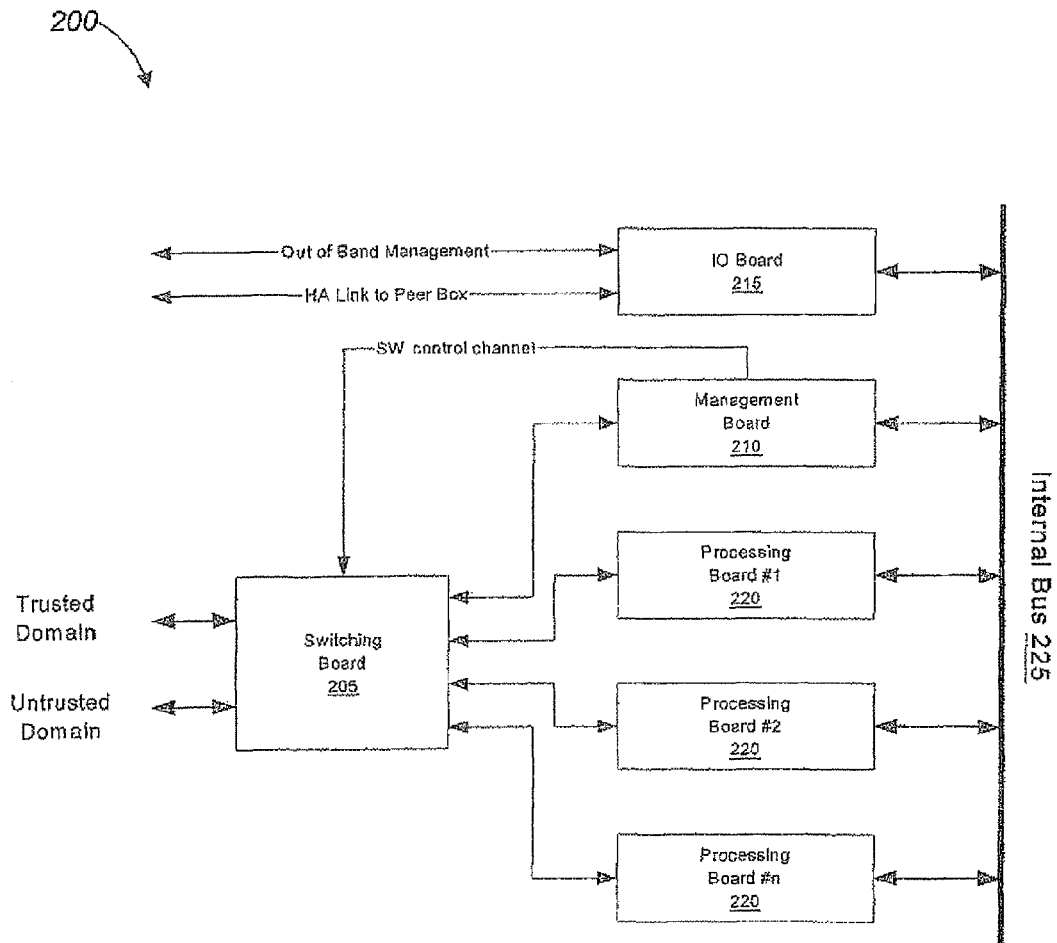


FIG. 2

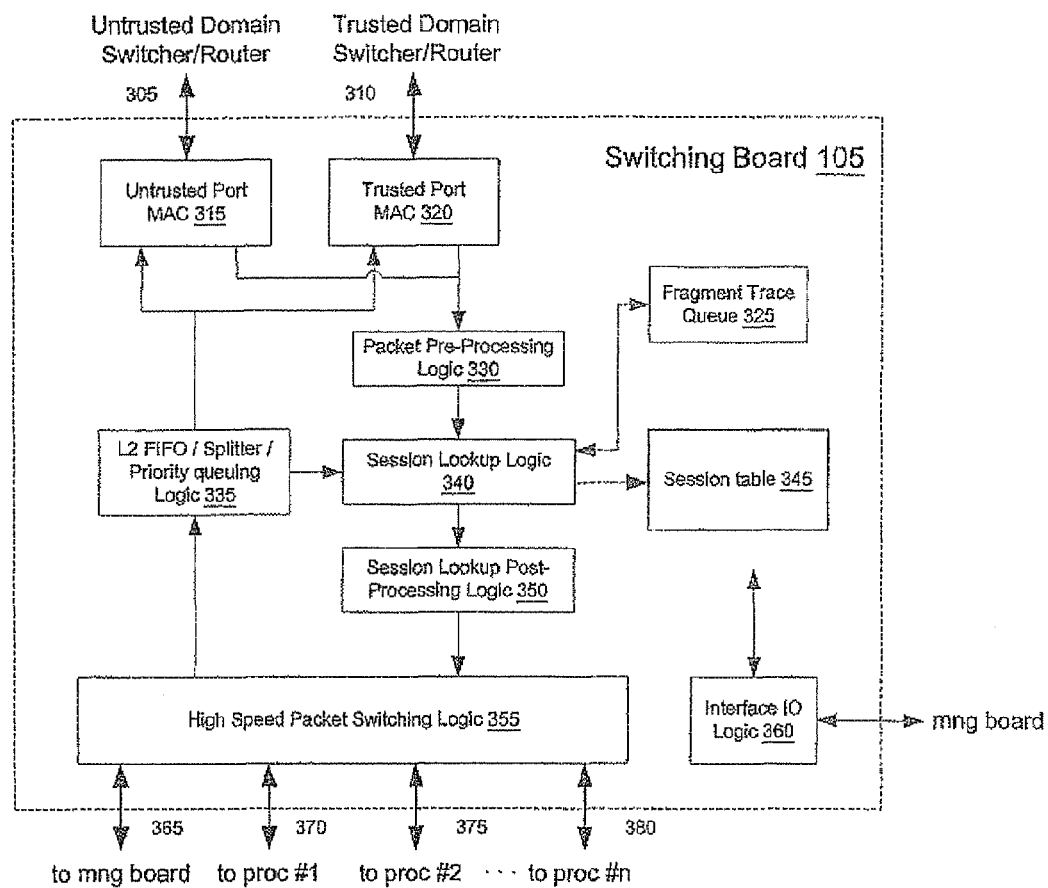


FIG. 3

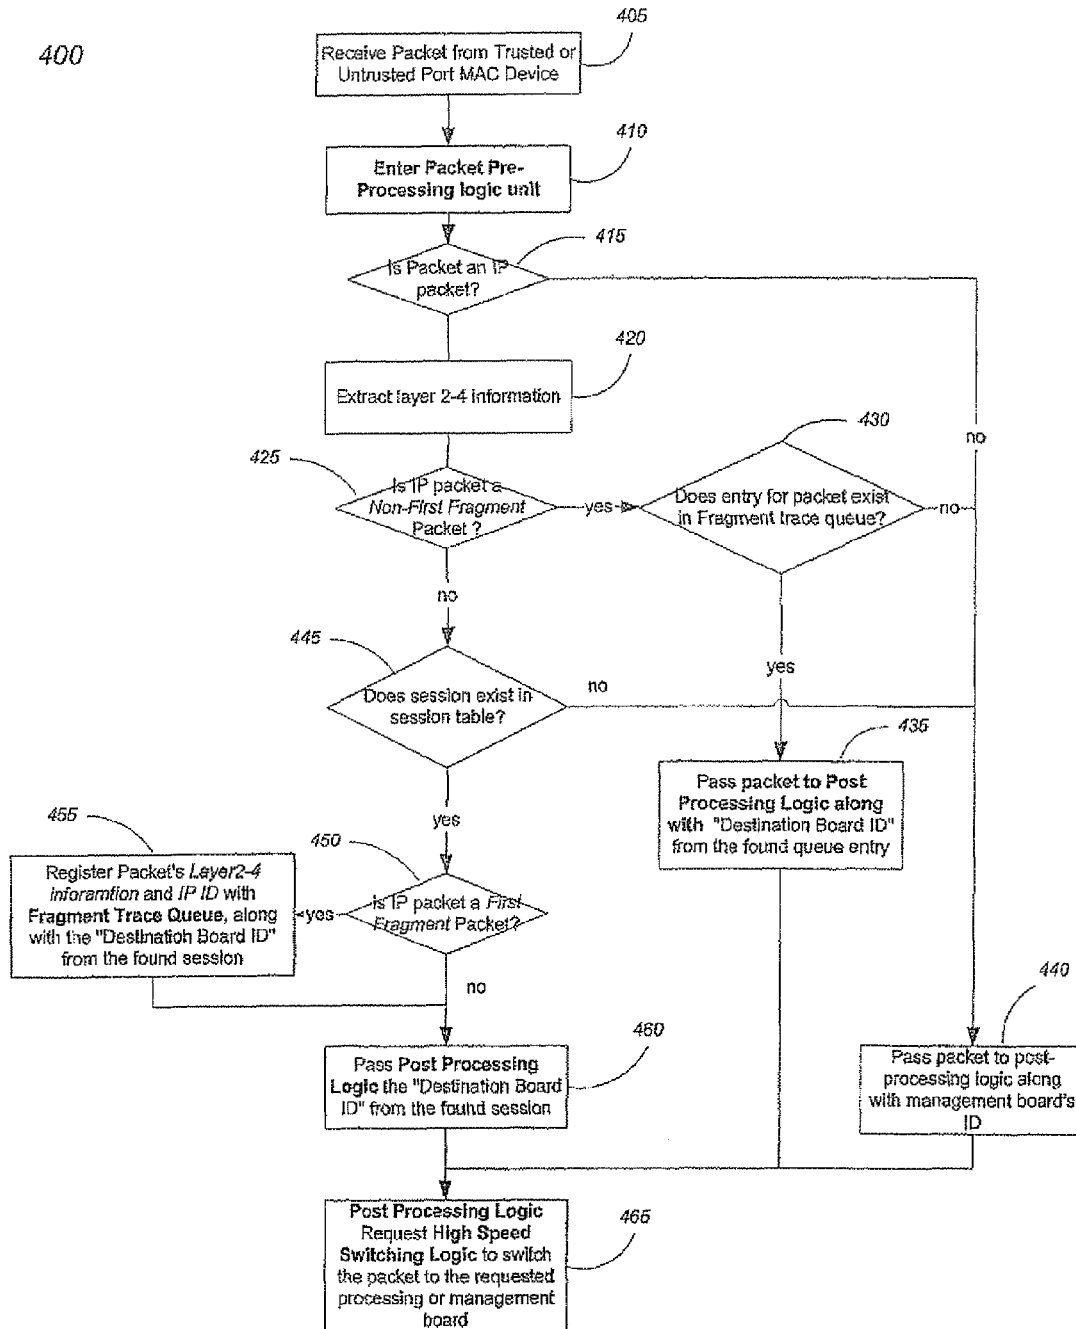


FIG. 4

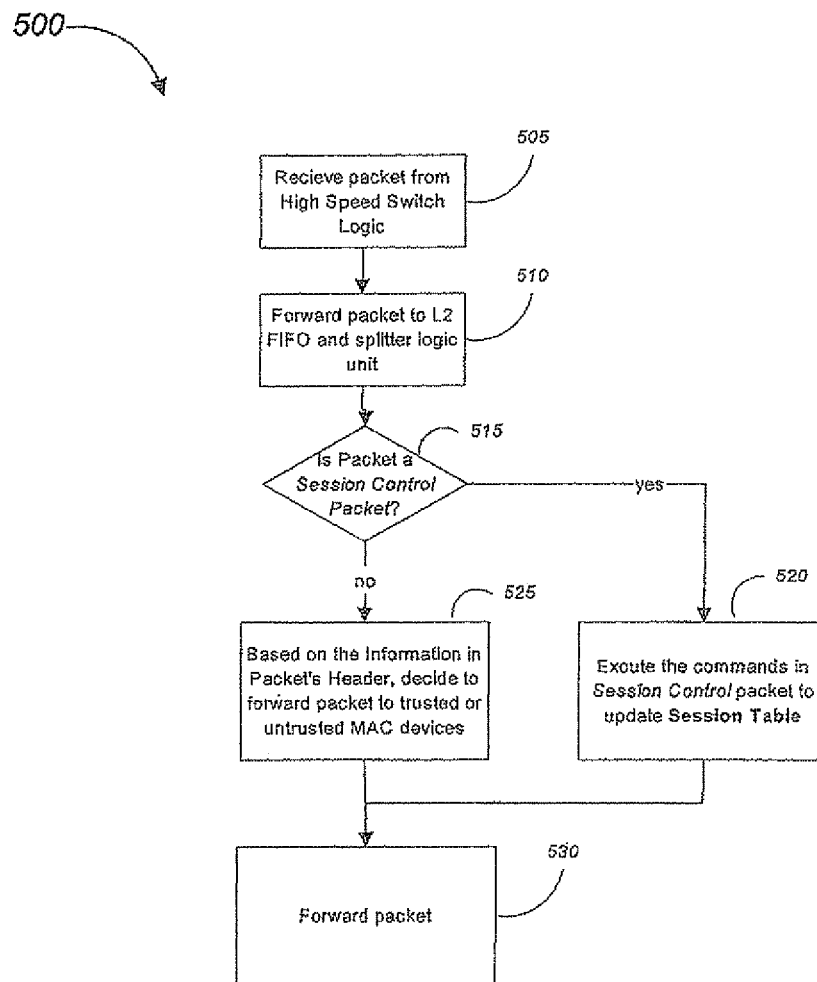


FIG. 5

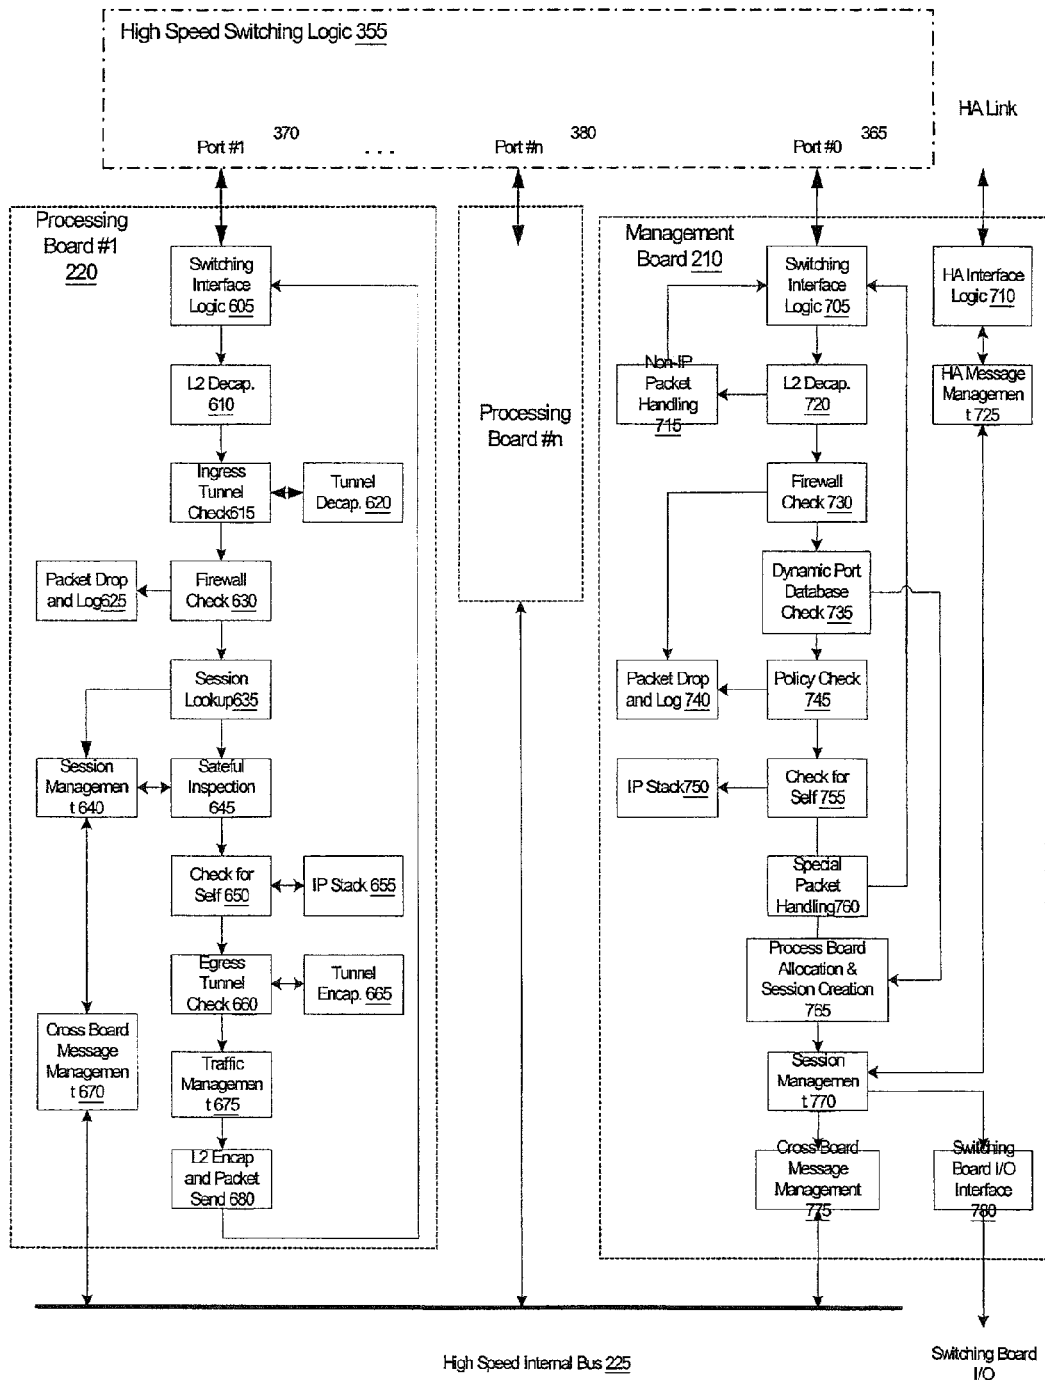


FIG. 6



**NETWORK SECURITY DEVICE****CROSS-REFERENCE TO RELATED APPLICATIONS**

This application is a continuation of U.S. patent application Ser. No. 13/302,808 filed Nov. 22, 2011 (now U.S. Pat. No. 8,654,779), which is a continuation of U.S. patent application Ser. No. 12/551,034 filed Aug. 31, 2009 (now U.S. Pat. No. 8,068,487), which is a continuation of U.S. patent application Ser. No. 11/428,235 filed Jun. 30, 2006 (now U.S. Pat. No. 7,602,775), which is a continuation of U.S. patent application Ser. No. 10/112,924 filed Mar. 28, 2002 (now U.S. Pat. No. 7,095,716), the entire contents of which are incorporated herein by reference. U.S. patent application Ser. No. 10/112,924 claims the benefit of prior U.S. provisional application 60/280,684, filed Mar. 30, 2001.

**BACKGROUND**

The invention relates to an Internet security system. There are many emerging trends in the communications world, including the increase in network technology and the proliferation of data networks. One example of a data network is a packet switch communication system. A conventional packet switch communication system includes one or more routers connecting a plurality of users. A packet is the fundamental unit of transfer in the packet switch communication system. A user can be an individual user terminal or another network. A router is a switching device which receives packets containing data or control information on one port, and based on destination information contained within the packet, routes the packet out another port to the destination (or intermediary destination). Conventional routers perform this switching function by evaluating header information contained within the packet in order to determine the proper output port for a particular packet.

The network can be an intranet, that is, a network connecting one or more private servers such as a local area network (LAN). Alternatively, the network can be a public network, such as the Internet, in which data packets are passed over untrusted communication links. The network configuration can include a combination of public and private networks. For example, two or more LAN's can be coupled together with individual terminals using a public network such as the Internet. When public and private networks are linked, data security issues arise. More specifically, conventional packet switch communication systems that include links between public and private networks typically include security measures for assuring data integrity.

To ensure security of communications, network designers have either incorporated security appliances, such as firewalls and traffic management devices in their systems or enhanced their routers with these functionalities. A firewall is an Internet security appliance designed to screen traffic coming into and out of a network location. A virtual private network provides a secure connection through a public network such as the Internet, between two or more distant network appliances using virtual private networking technology. In many cases, such appliances, although necessary from a security point of view, tend to act like bottlenecks in the network in terms of throughput.

**SUMMARY**

In general, in one aspect, this invention provides methods and apparatus, including computer program products, imple-

menting and using techniques for processing a data packet. An input port receives a data packet, a switching board classifies the data packet, determines whether the data packet should be accepted, and switches the data packet to a management board if the data packet is a first data packet in a session, and to a processing board if the data packet is not a first data packet in a session. A management board receives a data packet from the switching board, examines the data packet and forwards the data packet to one of the processing boards. One or more processing boards receives non-first data packets from the switching board and data packets from the management board and processes the data packets.

Advantageous implementations can include one or more of the following features. The classification of the data packet by the switching board can include determining if the data packet is an IP packet. The switching board can switch the data packet to the management board if the data packet is not an IP packet. If the data packet is an IP packet, the switching board can extract IP information from the data packet and determine if the data packet is a first data packet in a session. The IP information extracted from the data packet can include one or more of VLAN, IP source address, IP destination address, IP protocol number, IP identification number, and transport layer information. Switching the data packet based on the classification can include switching the data packet to a processing board, if the data packet is a non-first IP packet. The switching board can switch a non-first fragment IP packet to a processing board based on queue entries for a corresponding first fragment IP packet in a fragment trace queue on the switching board.

The management board can examine the data packet by performing one or more of the following checks on the data packet: a firewall check, a dynamic port database check, and an access policy check. The management board can forward the data packet by determining a current load on all processing boards and forwarding the data packet, based on the determined current load on all processing boards, to one of the processing boards. The data packet can be forwarded to the processing board having a lowest load. The management board can forward the packet to a processing board based on a user-selected criterion. The number of switching boards, management boards and processing boards can be configurable by a user. The checks performed on a data packet can be configurable by a user. The device can be an Internet security device. Two or more processing boards can be included that can process data packets in parallel. The one or more processing boards can process the received data packets by performing one or more of: policy enforcement, virtual private network traffic encryption and decryption, key negotiation, firewall protection, network address translation, tunnel encapsulation and de-capsulation, and hot backup between processing boards.

In general, in another aspect, this invention provides methods and apparatus, including computer program products and a firewall device, implementing and using techniques for processing a data packet. An input port receives a data packet. A switching board classifies the data packet, determines whether the data packet should be accepted by the firewall device, and switches the data packet to a management board if the data packet is a first data packet in a session, and to a processing board if the data packet is not a first data packet in a session. A management board receives a data packet from the switching board, performs a firewall check on the received data packet, examines the data packet, and forwards the data packet to one of the processing boards. One or more process-

ing boards receives non-first data packets from the switching board and data packets from the management board, and processes the data packets.

In general, in one aspect, this invention provides methods and apparatus, including computer program products and a secure gateway with firewall and virtual private network functionality, implementing and using techniques for processing a data packet. An input port receives a data packet. A switching board classifies the data packet, determines whether the data packet should be accepted by the gateway, and switches the data packet to a management board if the data packet is a first data packet in a session, and to a processing board if the data packet is not a first data packet in a session. A management board receives a data packet from the switching board, performs a firewall check on the received data packet, examines the data packet, and forwards the data packet to one of the processing boards. One or more processing boards receives non-first data packets from the switching board and data packets from the management board, and processes the data packets, the processing including virtual private network traffic encryption and decryption.

The invention can be implemented to realize one or more of the following advantages. A high-performance parallel processing system is provided for securing enterprise intranets, e-business operations and high-speed Internet access. The system includes a stateful-inspection firewall and VPN features and is built on a resilient platform that offers scalable resources targeting large enterprises and service providers. Parallel processing together with a fast firewall and encryption acceleration engine deliver high performance needed for broadband data applications. Service providers, such as Internet data centers (IDC), application infrastructure providers (AIP) and metropolitan area networks (MAN), can use the security systems to protect their infrastructures, and also as a basis for "in the cloud" managed security offerings. A scalable system architecture assures that customers are able to meet increasing bandwidth requirements using a single device. The system can be easily integrated into various types of environments with different requirements.

The details of one or more implementations of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

#### DESCRIPTION OF DRAWINGS

FIG. 1 is a schematic view of an environment in which a security device can be deployed.

FIG. 2 is a schematic view of a security device

FIG. 3 is a schematic view of a switching board of the security device of FIG. 2.

FIG. 4 is a flowchart showing how an incoming packet is processed by the switching board of FIG. 3.

FIG. 5 is a flowchart showing how an outgoing packet is processed by the switching board of FIG. 3.

FIG. 6 is a schematic view of the architecture and operation of a processing board and management boards.

Like reference symbols in the various drawings indicate like elements.

#### DETAILED DESCRIPTION

An Internet security device (200) in accordance with the invention can, for example, be applied in a computer network environment (100) such as the one shown in FIG. 1. A number of users (110, 115, 120) are connected to a wide area network

(WAN) (105), such as the Internet. The users may be individuals or different types of entities, such as corporations or universities. Some users (110) can be connected directly to the wide area network (105), while others are connected through some subnet, such as a local area network (LAN) (125) or a virtual network in which an administrator has defined a group of users as a single subnet. Information travels between the different users, entities, and groups as packets. An Internet security device (200) can be used to control what traffic passes between networks or users, for example, between the WAN (105) and an individual user (115) or between the WAN (105) and the LAN (125) to which one or more users (120) are connected. The Internet security device (200) can be configured to either protect the single user (115) or an entire network, such as LAN (125) and the users (120) who are connected to the LAN (125), or a defined user group operating on the LAN (125). Protection includes inspecting, screening and monitoring incoming and outgoing packets, including detecting, intercepting and dropping unwanted packets in accordance with defined policies. The policies can be pre-set by the user (115) or a network administrator. The Internet security device (200) will now be described in further detail, by way of example.

As can be seen in FIG. 2, an Internet security device (200) in accordance with one implementation of the invention includes a switching board (205), a management board (210), an I/O board (215), and several processing boards (220). The boards are installed in a single chassis box and communicate with each other through one or more internal buses (225). The security device (200) is modular and additional resources can be added to increase performance. The different boards and the function of each board in the security device (200) will now be described in further detail.

#### The Switching Board (205)

FIG. 3 shows a more detailed schematic view of the switching board (205). The switching board (205) receives and processes incoming packets from the WAN (105) or LAN (125) and determines for each packet if the packet should be forwarded to the management board (210) or to a processing board (220) in the security device (200). The switching board (105) includes the following components. A trusted domain switcher/router port (310) and an untrusted domain switcher/router port (305), are used to receive packets. A packet pre-processing logic unit (330) is used to examine if a packet is an IP packet. A post-processing logic unit (350) is used to switch packets to a management board port (365) on a high-speed packet switching logic unit (355). A session lookup logic unit (340) is used to determine if a packet is a non-first fragment packet using a session table (345). A fragment trace queue (325) contains queue entries relating to packets that have passed through the Internet security device, thereby enabling non-first fragment IP packets to pass straight to a proper processing board (370, 375, 380). An interface I/O logic component (360) handles the communication between the management board (210) and the switching board (105). A layer 2 FIFO, splitter and priority queuing logic unit (335) is used to buffer packets, determine if a packet is a session control packet, and queue packets in a priority queue so that certain packets get priority over others. The priority mechanisms for the priority queue can be configured by a system user according to any criterion he or she may select. For example, control packets can have priority over normal packets. The priority queuing mechanism ensures quality of service and high throughput.

FIG. 4 shows a schematic flowchart of a process (400) describing how an incoming packet is routed through the switching board (205) to a processing board (370, 375, 380)

5

or to a management board (365). The process starts by receiving a packet from the LAN (125) or WAN (105) in FIG. 1. The packet is received using a trusted port media access connection (MAC) (320) or an untrusted port MAC (330) from either the trusted domain switcher/router port (310) or from the untrusted domain switcher/router port (305), respectively (step 405). The packet first enters the packet pre-processing logic unit (330) (step 410). The packet pre-processing logic unit examines if the packet is an IP packet (step 415). If the packet is not an IP packet, then the packet is forwarded to the post-processing logic unit (350) (step 440), which switches the non-IP packet to the management board port (325) on the high-speed packet switching logic unit (355) (step 465), which ends the process.

If the packet pre-processing logic unit (330) in step 415 determines that the packet is an IP packet, all the layer information is extracted from the data packet (step 420). The layer information in the packet includes a VLAN ID, if present, in layer 2; IP source address and destination address, IP protocol number and IP ID in layer 3; and transport layer identification information, such as TCP/UDP source port and destination port, ESP/AH SPI number, and PPTP/L2TP call ID in layer 4.

The packet then passes to the session lookup logic unit (340), which determines if the header of the IP packet indicates that the packet is a non-first fragment packet (step 425). If the packet is a non-first fragment packet, the packet's layer information (in one implementation the IP source address, IP destination address and IP ID, along with the VLAN ID, if present) is used to check if a corresponding queue entry exists in the fragment trace queue (325) (step 430).

If a queue entry is found, then the packet is sent through the post-processing logic unit (350) and high speed packet switching logic unit (355) to the port indicated by the queue entry (step 435). If no queue entry is found, the packet is sent to the management board port (325) (step 440).

Returning now to step 425, if the process decides that the IP packet is not a non-first fragment packet, that is, the packet is a first fragment IP packet, the packet is sent to the session lookup logic unit (340), where the process searches a session recorder in a session table (345) to find an associated session (step 445). In one implementation, the information used to search the session table is the extracted layer 2-4 information, that is, VLAN ID, source IP, destination IP, protocol number, source port and destination port. If no session is found, then the packet is forwarded to the management board port (325) (step 440).

If there is a session that matches the packet information, the process checks if the packet is a first fragment of a fragmented IP packet (step 450). If so, the process registers the packet's layer 2-4 information (that is, source IP, destination IP, protocol number) and IP ID with the fragment trace queue (325) along with the target port number found in the session entry (step 455), so that any subsequent non-first fragment IP packet can pass straight to the correct processing board.

If it is decided in step 450 that the packet is not a first fragment packet, the packet is sent through the session lookup post-processing logic unit (350). Here the destination board ID from the found session is used (step 460) to instruct the high-speed packet switching logic unit (355) to forward the packet to the port number indicated in the matching session's target port field. Finally, the high-speed packet switching logic unit (355) switches the packet to the requested processing board (step 465). This completes the processing of incoming packets on the switching board (105).

FIG. 5 is a flowchart showing a process (500) for how an outgoing packet is processed by the switching board (105) in FIG. 3, that is, a packet coming into the switching board (105)

6

from one of the processing board ports (330-340) that is destined for the entrusted domain switch/router (305) or the trusted domain switch/router (310). The process begins with the receipt of the outgoing packet by the high-speed packet switching logic unit (355) from one of the processing boards (120) through one of the processing ports (330-340) (step 505).

The outgoing packet from the high-speed packet switching logic unit 355 is forwarded to the L2 FIFO, splitter and priority queuing logic unit (335) (step 510). The L2 FIFO, splitter and priority queuing logic unit (335) buffers the packet and determines if the packet is a session control packet (step 515). In one implementation, the splitter classifies different packet types based on the Ethernet type field in the layer 2 Ethernet header. In one implementation, the received packet is a session control packet if the packet carries a special MAC type value. The process then sends the session control packet to the session lookup table (340), where commands in the session control packet are executed and the entries in the session table are updated (step 520).

If the packet is not a session control packet, then the splitter decides to which MAC port the packet should be forwarded (step 525). In one implementation, the splitter looks at the packet's source MAC address inside the layer 2 Ethernet header to decide to which MAC port the packet should be forwarded. Finally the packet is forwarded to the determined location (step 530), which ends the process.

The Management Board

As was described above, the management board (210) receives and processes the first packet in a session, as well as any packets that cannot be directly switched to any of the processing boards (220). The handling of the first packet on the management board (210) can include packet classification, policy lookup, session setup (including VPN sessions), load balancing (the session loads to different processing boards), synchronizing session state information with the switching board and high availability (HA) peers, and handling management tasks. All of these tasks will be described in further detail below.

The processing board (220) and the management board (210) are schematically depicted in FIG. 6 as part of a high-level system architecture that shows the interconnection of the boards. The management board (210) contains the following components. A switching interface logic unit (705) is connected to the High Speed Packet Switching Logic unit on the switching board (205) and receives packets to be processed.

Packet classification includes the extraction of layer information and "typing" of a packet. A layer 2 de-capsulation unit (720) extracts layer information about a packet from a layer 2 portion of the packet header. A non-IP Packet Handling unit (715) handles packets that are not IP packets. How non-IP packets are handled is discussed in greater detail below.

The policy lookup operations can include a firewall inspection including the enforcement of global and local policies. A Firewall Check engine (730) checks packets for attacks configured by a system user. A log (740) contains information about possible attacks. An Access Policy check unit (745) can be initiated for access policies created by a firewall administrator. A special packet handling unit (760) performs various types of checks on packets, for example, whether a packet's TTL (time to live) value has reached zero.

Session set-up includes creation of new sessions for each first packet received that is identified to be part of a session. Certain live session applications may create other, dynamic sessions, for example, an FTP port command can require a session to be created. A dynamic port database check unit

(735) contains dynamic session information for dynamic sessions that is created by such live session applications. A new session creation unit (765) is used to create new sessions. Session state information is synchronized with the switching board and other high availability (HA) peers to speed the processing of packets through the system. A cross board message management unit (775) is used to send session creation messages to a selected processing board. The cross board message management unit (670) on the target processing board is responsible for notifying the session management unit (640) on the processing board (220) of a session creation, so the session management unit (640) in turn can create a local session record on the target processing board (220). A switching session table (345) on the switching board (205) is used by a switching board I/O interface (780) on the management board (210) to notify the switching board (205) about a session creation and a new home processing board number for the session. A session management unit (770) removes a session entry from the session database, creates a log of sessions that have finished, and notifies the switching session table (345) of the removal.

A packet testing unit (755) is used to see if a packet is addressed to the system itself. An IP stack (750) on the management board (210) contains information for packet processing.

As can be seen from FIG. 6, the flow of data in the management board (210) starts with a data packet entering the switching interface logic unit (705) from the high speed switching logic port #0 (325) that is connected to the High Speed Packet Switching Logic unit on the switching board. (205)

The data packet continues to the layer 2 de-capsulation unit (720) where layer information about the packet is extracted from layer 2. Layer 2 contains the ULAN ID information. The extracted layer 2 information is used to check if the packet is designated for this system. If the layer 2 information shows that the packet is not designated for the system, the packet is to be dropped and the data packet flow for this packet ends. If the layer 2 information shows that the packet is designated for the system, the management board (210) checks the layer 3 information (that is, the IP source address, the IP destination address, and the IP protocol number) to see if the packet is an IP packet. If the packet is a non-IP packet, the packet is to be sent to the NON-IP Packet Handling unit (715) for further processing. Such processing includes ARP (address resolution protocol) handling, NSRP (Netscreen Redundancy Protocol) messaging, and so on.

If the packet is an IP packet, the packet is sent to a Firewall Check Engine (730) that checks the packet. In one implementation, the packet is screened for attacks configured by a system user. The checks can include, but are not limited to, detection of so-called "Port Scan," and "Ping of Death," attacks. If any of these attack types are detected, the packet is interpreted to be an attempted attack to the system and is therefore dropped and logged together with the attack information (740).

Some of the layer 2, layer 3 and layer 4 information can be carried along with the packet as the packet moves among the units on the management board (220). This information includes: VLAN ID from layer 2, if a VLAN ID is present; source address, destination address, and protocol number, from layer 3; and TCP/UDP source port number, destination port number, from layer 4. The information can uniquely identify a network traffic connection (session).

After the firewall check is completed, the data packet passes through a dynamic port database check (735). The dynamic port database contains the dynamic session infor-

mation that is created by certain live session applications. One example of such an application is an FTP application. When an FTP connection starts (FTP control session), an FTP data tunnel session may be dynamically initiated from the FTP server by a command in the FTP control session. In order to capture such session creation, an entry must be dynamically created in the database once such command is encountered on the FTP control session. This allows a data session initiated from the FTP sever to be matched once it arrives and the system can correctly create a new session (765). Once a session matches a record, that matching record is removed from the dynamic database. If no session matches are detected for a given record (that is, over a predetermined time span), the record is automatically removed by system as well.

If no match is found in the dynamic database, an Access Policy check (745) is initiated for any access policies created by a firewall administrator. The policies typically contain rules about what kinds of traffic should be allowed and how the traffic should be managed. For example, a policy can state "From Host A to web server B, all HTTP traffic is allowed." If there is no policy matching the session information, or the policy indicates a deny action, the packet is dropped and logged. Otherwise, the packet is tested (755) to see if the packet is addressed to the system itself. Such traffic includes, for example, network management packets. If the packet is addressed to the system itself, the packet is copied to the IP stack (750) on the management board (110) for further processing. If the packet is not addressed to the system itself, a check for special packet handling (760) is performed, for example, to check whether the packet's TTL (time to live) value has reached zero. A session is then created using policy information to enforce the stateful inspection for all the packets that belong to the session. (765).

The session creation includes three steps: In a first step, a session record is created in the management board session database (765). In a second step, one of the available processing boards is selected to host the session. In one implementation, a processing board (220) is selected based on the current load on the respective processing boards. In one implementation, a processing board (220) is selected in such a way that the load be distributed as evenly as possible across all the processing boards. After a processing board has been selected, a session creation message is issued along with all fatal session information and the data packet itself. The session creation message is then sent to the selected processing board using the cross board message management unit (775). Each processing board (220) includes a cross board message management unit (670). The cross board message management unit (670) on the target processing board (220) is responsible for notifying the session management unit (640) on the processing board of a session creation, so the session management unit (640) in turn can create a local session record on the target processing board (220).

In a third step, the session management unit (770) sends a command to the switching session table (345) on the switching board (205), using the switching board 110 interface (780) to notify the switching board about the session creation and the new home processing board number for the session. Thereby, the following packets on the session can be correctly directed to the selected processing board (220).

Where there is a session creation, there is also a session removal once the session lifetime is up. The session removal action is initiated from the selected processing board (220). Once the stateful inspection engine (645) on a processing board (220) detects a session close, the stateful inspection engine (645) notifies the session management unit (640), which creates a delete session command and sends the delete

session command to the session management unit (770) on the management board (210) through the cross board message management units (670, 775). Upon receiving such commands, the session management unit (770) removes the session entry from the session database, creates a log of session finished and notifies the switching session table (345) on the switching board (205) of the removal. This completes the description of the data flow for the data packet on the management board (210).

#### The Processing Board

Turning now to the processing boards (220), as was mentioned above, the Internet security device (200) may include several processing boards (220). The scalability aspect of the Internet security device (200) makes it possible for a user or administrator of the Internet security device (200) to add as many processing boards as needed in order to achieve desired performance and throughput. Each of the processing boards (220) receives packets from the switching board (205). The packets are part of a session that has been identified and created in accordance with the management board (210). The processing boards (220) handle tasks such as policy enforcement, VPN (Virtual Private Network) traffic encryption and decryption, key negotiation, firewall protection, NAT (Network Address Translation) and hot backup between boards. These tasks will be described in further detail below.

In order to simplify the description, only one processing board flow will be described here by way of example. Each processing board (220) has the same characteristics in the described implementation of the invention. Further, it is assumed that the first packet in a session has been processed by the management board (210), so that the switching board (205) has been made aware of where (i.e., to which processing board (220)) to send all subsequent packets in a given session.

In one implementation, all the processing boards (220) are identical to each other, from both physical and firmware points of view. Since the processing boards (220) have the same configuration, if one processing board fails, another board (or combination of boards) can easily take over the tasks of the failed processing board (220). However, during normal operation, each processing board (220) is responsible for processing only the sessions that have been assigned by the management board (210). If a processing board (220) receives an "unrecognized packet," the processing board (220) returns the packet to the management board (210) for further investigation.

In one implementation, each processing board (220) contains the following components. A high speed switching logic port#1 (315) receives packets from a switching interface logic unit (605). A layer 2 de-capsulation unit (610) extracts layer 2 information to check if a packet is designated for the system. A traffic management unit (675) enforces a matching traffic management policy configured by a system administrator. A layer 2 encapsulation and send unit (680) is used to encapsulate and send a packet out from the processing board (220).

An ingress tunnel check unit (615) is used to find a proper tunnel for a packet. A decode unit (620) is used to call the de-capsulation vector from the tunnel to decode encrypted octets back to a plain IP packet (620). An egress tunnel check unit (660) checks if the packet is addressed to an egress tunnel. A tunnel encapsulation unit (665) encrypts data packets so that packets can reach a target through a defined tunnel.

An IP firewall attack check unit (630) is used to check packets for attacks or illegal information. A packet information log (625) is used to log dropped packets. A local session table (635) is used to keep track of sessions handled by the processing board (220). A check for self unit (650) is used to check if a packet is addressed to the system itself.

A stateful inspection engine (645) tracks the session state of a packet; monitors dynamic session creation and closing and other irregularities of the packet flow, at high network layers, such as application layers, and so on. A session management unit (640) sends messages to the management board session management unit through a cross board message management unit (670, 775).

The data flow in processing board #1 (220) starts when a data packet enters the switching interface logic unit (605) from the high speed switching logic port#1 (315). Just like in the management board data flow, the packet passes through a layer 2 de-capsulation unit (610) where the layer 2 information is extracted to check if the packet is designated for the system. If the packet is designated for the system, all the necessary layer 2, layer 3 and layer 4 information is extracted out of the data packet and carried along with the packet throughout the packet processing. In one implementation, the extracted information includes: VLAN ID from layer 2, if it is present; IP Source Address, Destination Address and Protocol Number from layer 3; TCP/UDP source port number and destination port number, ESP/AH protocol SPI number, and GRE protocol call ID from layer 4.

The data packet received by the processing board can be encapsulated by a tunnel. One example is IPsec encrypted VPN traffic. When VPN tunnels are used, only the VPN tunnel for which the packets are intended can decrypt the packets and recover the packets to obtain the meaningful data packets. The Internet security device (200) builds all the mechanisms required to support the VPN tunnel. Once a packet comes in to the processing board (220), an ingress tunnel check unit (615) finds the proper tunnel for the packet, and calls the de-capsulation vector from the tunnel to decode the encrypted octets back to a plain IP packet (620). Once the packet has been decoded, the IP, TCP/UDP layer 3 and layer 4 data is updated to regenerate the IP packet for transmission to its intended recipient.

Once all the tunnel encapsulations (if there are any) have been removed, an IP firewall attack check (630) can be applied to the packet. If any attack or untrusted information is found in the packet information, then the packet is dropped and logged as irregular (625).

Next, the system discovers the session context for the incoming packet by performing a look-up operation using the local session table (635) of the processing board (220). As was mentioned in the management board data flow description above, the session context is created by commands from the management board (210) when the management board (210) encounters a new session stream in the management board's data flow. In most cases, the session context is found successfully, but in a case where no session can be found, the packet is sent to the management board (210) for further investigation (640, 670).

Once the system has the session context, stateful inspection can be conducted on the packet (645). The stateful inspection can include tracking the session state of the packet; monitoring dynamic session creation and closing and other irregularities of the packet flow, at high network layers, such as application layers, and so on. The stateful inspection engine (645) keeps the session state information updated while the packet flows through the processing board (220). If the stateful inspection engine (645) notices a session closing, the stateful inspection engine (645) notifies the session management unit (640), which in turn sends a message to the management board session management unit (770) through the cross board message management unit (670, 775), so that the session database can be updated, and the session in the switching board session table (345) can be freed up. Also, if

11

the stateful inspection engine (645) sees a dynamic session creation event, the stateful inspection engine (645) sends a “dynamic session register” command to the dynamic port database of the management board, so that the future spawned new session traffic can be captured by the dynamic port data-  
base checks. 5

After the packet passes the inspection provided by the stateful inspection engine (645) (e.g., the packet is authenticated), the packet enters a stage for egress processing. First, the packet is checked to determine if the packet is targeted to the stack of the system (650). If so, the packet is forward to the IP stack (655) where the flow ends. Examples for such traffic include IKE (Internet Key Exchange) key negotiation traffic, management traffic, and so on. 10

The packet then proceeds to the packet manipulation processing, where NAT/PAT (Network Address Translation/Port and Address Translation) is conducted if necessary as indicated by the session context. A NAT/PAT is an operation to replace at least some of the layer 3/layer 4 information with synthesized values, so that a portion of the original packet’s information (for example, identities) can be protected. 15 20

Next, the packet’s outgoing route is determined. For normal traffic, the determination merely involves route table look up and finding the right default gateway for the packet; for other traffic, if there is an egress tunnel associated with the session, the system calls the corresponding tunnel’s encapsulation vector (665) to encrypt the data packet so that the packet can reach the target through the defined tunnel. One such example is out going VPN traffic, in which the packet body needs to be encrypted before the packet enters the WAN (105) or LAN (125). 25 30

The packet then passes to a traffic management unit (675) where, as appropriate, a matching traffic management policy configured by the system administrator is enforced. Such a policy places the traffic to the proper priority and bandwidth pool, and also limits the maximum bandwidth the packet can consume. 35

In the last step, the system encapsulates the packet (680) (for example, to an Ethernet packet) and forwards the packet to the transmit interface of the switching interface logic unit (605) to send the packet back to the high speed switching logic unit (355), which in turn forwards the packet to the switch or router connected to the trusted or untrusted domain ports. This completes the description of the processing board (220). 40 45

A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, two or more switching boards can be provided, where each switching board processes a fraction of the total number of packets. More than one management board can be provided if the state information is shared between the different management boards. Accordingly, other embodiments are within the scope of the following claims. 50 55

What is claimed is:

1. A system comprising:  
a network device to:

- determine whether a data packet is a first data packet in a session associated with the data packet;
- determine whether information regarding the session, associated with the data packet, is stored in a data structure after determining whether the data packet is the first data packet in the session, the data structure storing information regarding sessions;

12

selectively transmit the data packet to one or more first components, of the network device, or to one or more second components of the network device,

the one or more second components being different than the one or more first components,

the data packet being transmitted to the one or more first components when the information regarding the session, associated with the data packet, is not stored in the data structure, or

the data packet being transmitted to the one or more second components, without being transmitted to the one or more first components, when the information regarding the session,

associated with the data packet, is stored in the data structure; and

selectively process the data packet by one of:

the one or more first components, or

the one or more second components.

2. The system of claim 1, where, when determining whether the data packet is the first data packet in the session, the network device is to:

extract information from the data packet, and

determine whether the data packet is the first data packet based on the information extracted from the data packet.

3. The system of claim 1, where, when selectively processing the data packet, the network device is to:

monitor, using the one or more second components, irregularities of a packet flow associated with the data packet.

4. The system of claim 1, where, when selectively processing the data packet, the network device is to:

examine the data packet using the one or more first components,

when examining the data packet, the one or more first components are to perform at least one of:

a firewall check on the data packet, or

an access policy check on the data packet.

5. The system of claim 4, where the network device is to perform the access policy check,

where, when performing the access policy check, the network device is to:

determine whether a policy associated with the session exists,

determine whether the policy authorizes packets, associated with the session, to be transmitted to a destination of the packets, and

cause the data packet to be transmitted to a destination of the data packet when the policy authorizes the packets to be transmitted.

6. The system of claim 4, where the network device is to perform the firewall check,

where, when performing the firewall check, the network device is to:

determine whether the data packet is associated with an attack, and

when the data packet is associated with the attack:

drop the data packet, and

log information regarding the attack.

7. The system of claim 4, where the network device is further to:

determine, using the one or more first components, a load on each processing board of a plurality of processing boards of the network device, and

forward, using the one or more first components, the data packet to a particular processing board of the plurality of processing boards after examining the data packet,

## 13

the load on the particular processing board being a lowest load out of the load of each processing board of the plurality of processing boards, and the particular processing board including the one or more second components.

8. The system of claim 4, where, when examining the data packet, the one or more first components are to:

perform the firewall check,  
perform a dynamic port database check after performing the firewall check, and  
perform the access policy check after performing the dynamic port database check.

9. A method comprising:

determining, by a network device, whether a data packet is a first data packet in a session associated with the data packet;

determining, by the network device, whether information regarding the session, associated with the data packet, is stored in a data structure based on determining whether the data packet is the first data packet in the session, the data structure storing information regarding sessions;

selectively transmitting, by the network device, the data packet to one or more first components of the network device, or to one or more second components of the network device,

the one or more second components being different than the one or more first components, the data packet being transmitted to the one or more first components when the information regarding the session, associated with the data packet, is not stored in the data structure; or

the data packet being transmitted to the one or more second components, without being transmitted to the one or more first components, when the information regarding the session, associated with the data packet, is stored in the data structure; and

selectively processing, by the network device, the data packet by one of:  
the one or more first components, or  
the one or more second components.

10. The method of claim 9, where selectively processing the data packet includes:

examining the data packet using the one or more first components,  
where examining the data packet includes performing at least one of:  
a firewall check,  
a dynamic port database check, or  
an access policy check.

11. The method of claim 10, where examining the data packet includes:

performing the firewall check,  
performing the dynamic port database check after performing the firewall check, and  
performing the access policy check after performing the dynamic port database check.

12. The method of claim 10, where examining the data packet includes performing the firewall check, and where performing the firewall check includes:

determining whether the data packet is associated with an attack, and  
when the data packet is associated with the attack:  
dropping the data packet, and  
logging information regarding the data packet and information regarding the attack.

## 14

13. The method of claim 10, where examining the data packet includes performing the access policy check, and where performing the access policy check includes:

determining whether a policy associated with the session exists,

determining whether the policy authorizes packets, associated with the session, to be transmitted to a destination of the packets, and

causing the data packet to be transmitted to a destination of the data packet when the policy authorizes the packets to be transmitted.

14. The method of claim 9, further comprising:

examining, using the one or more first components, the data packet;

determining, using the one or more first components, a load on each processing board of a plurality of processing boards of the network device; and

forwarding, using the one or more first components, the data packet to a particular processing board of the plurality of processing boards after examining the data packet,

the load on the particular processing board being a lowest out load of the load of each processing board of the plurality of processing boards.

15. The method of claim 9, where selectively processing the data packet includes:

monitoring, using the one or more second components, irregularities of a packet flow associated with the data packet.

16. A non-transitory computer-readable medium storing instructions, the instructions comprising:

one or more instructions which, when executed by a device, cause the device to determine whether a data packet is a first data packet in a session associated with the data packet;

one or more instructions which, when executed by the device, cause the device to determine whether information regarding the session, associated with the data packet, is stored in a data structure after determining whether the data packet is the first data packet in the session, the data structure storing information regarding sessions;

one or more instructions which, when executed by the device, cause the device to selectively transmit the data packet to one or more first components, of the device, or to one or more second components of the device, the one or more second components being different than the one or more first components,

the data packet being transmitted to the one or more first components when the information regarding the session, associated with the data packet, is not stored in the data structure, or

the data packet being transmitted to the one or more second components, without being transmitted to the one or more first components, when the information regarding the session, associated with the data packet, is stored in the data structure; and

one or more instructions which, when executed by the device, cause the device to selectively process the data packet by one of:

the one or more first components, or  
the one or more second components.

17. The non-transitory computer-readable medium of claim 16, where the one or more instructions to selectively process the data packet include:

15

one or more instructions to perform, using the one or more first components, at least one of:  
a firewall check, or  
an access policy check.

18. The non-transitory computer-readable medium of claim 17, where the one or more instructions to perform the at least one of a firewall check or an access policy check include:  
one or more instructions to perform the firewall check,  
one or more instructions to perform a dynamic port database check after performing the firewall check, and  
one or more instructions to perform the access policy check after performing the dynamic port database check.

19. The non-transitory computer-readable medium of claim 16, where the one or more instructions to selectively process the data packet include:  
one or more instructions to monitor, using the one or more second components, irregularities of a packet flow associated with the data packet.

20. The non-transitory computer-readable medium of claim 16,  
where the one or more instructions to determine whether the data packet is the first data packet in the session include:  
one or more instructions to extract information from the data packet, and  
one or more instructions to determine whether the data packet is the first data packet based on the information extracted from the data packet.

\* \* \* \* \*

16